

ARTICLE

Received 24 Sep 2016 | Accepted 29 Nov 2016 | Published 16 Jan 2017

DOI: 10.1038/ncomms14106

OPEN

Prediction and real-time compensation of qubit decoherence via machine learning

Sandeep Mavadia^{1,2,*}, Virginia Frey^{1,2,*}, Jarrah Sastrawan^{1,2}, Stephen Dona^{1,2} & Michael J. Biercuk^{1,2}

The wide-ranging adoption of quantum technologies requires practical, high-performance advances in our ability to maintain quantum coherence while facing the challenge of state collapse under measurement. Here we use techniques from control theory and machine learning to predict the future evolution of a qubit's state; we deploy this information to suppress stochastic, semiclassical decoherence, even when access to measurements is limited. First, we implement a time-division multiplexed approach, interleaving measurement periods with periods of unsupervised but stabilised operation during which qubits are available, for example, in quantum information experiments. Second, we employ predictive feedback during sequential but time delayed measurements to reduce the Dick effect as encountered in passive frequency standards. Both experiments demonstrate significant improvements in qubit-phase stability over 'traditional' measurement-based feedback approaches by exploiting time domain correlations in the noise processes. This technique requires no additional hardware and is applicable to all two-level quantum systems where projective measurements are possible.

¹ARC Centre for Engineered Quantum Systems, School of Physics, The University of Sydney, Sydney, New South Wales 2006, Australia. ²National Measurement Institute, West Lindfield, New South Wales 2070, Australia. * These two authors contributed equally to this work. Correspondence and requests for materials should be addressed to M.J.B. (email: michael.biercuk@sydney.edu.au).

The applications of quantum-enabled technologies are compelling and already demonstrating significant impacts, especially in the realm of sensing^{1–5} and metrology⁶. However, in nearly all applications the phenomenon of decoherence—effectively the randomization of a quantum system's state by the environment—limits the viability of quantum technologies. In the case of qubits, fundamental building blocks in many applications, the net result is that the useful lifetime of the qubit state is shortened, reducing their deployability for quantum information⁷, quantum simulation^{8–13} or other applications. Methodologies for stabilising qubits against decoherence represent a critical need in quantum technology.

Control engineering¹⁴ techniques are emerging as a promising alternative to engineering passive robustness at the device level in realising stable quantum systems^{15–18}. Beyond widely adopted open-loop control^{18–20}, a qubit subjected to stochastic evolution of its phase degree of freedom—dephasing (inset Fig. 1a)—can be stabilised by cyclically performing measurements on the qubit and then compensating for the measured phase evolution in a feedback loop^{21–23}. However, so far, feedback control^{24–28} has largely been limited by state-collapse under projective measurement, mandating access to weak measurements²² or ancilla states²⁹, or largely sacrificing useful quantum coherence in the controlled system²³.

Our objective is to enhance the performance of incoherent feedback stabilization (that is, using only classical information) of a qubit experiencing dephasing while also relaxing the need for projective measurements. Our approach is based on predictive control; a variety of techniques in filtering^{14,30–32} and machine learning³³ allow the estimation of future state evolution based on past measurement outcomes of the system. Here, we deploy a well established algorithm from machine learning to learn about a random dephasing process affecting a qubit, and then predict the impact of future dephasing based only on standard projective measurements. We use this information to perform real-time stabilization of the qubit state during periods in which the qubit is unsupervised but still subject to stochastic dephasing. Our method exploits the presence of commonly encountered temporal correlations in the dephasing process³⁴ to allow future prediction; no deterministic model of qubit state evolution is required. To the best of our knowledge, despite its ubiquity in classical settings, predictive control has not been employed in the context of quantum-coherent technologies.

Results

Supervised learning based on qubit-phase measurements. In the language of machine learning, we consider the qubit's instantaneous phase which we would like to predict at a future discretized time, t_k , as labels, $\phi^P(t_k)$, and an arbitrary number, n , of previous measurements, ϕ_i^M (indexed by i and obtained by any appropriate method), as their associated features. We then calculate a linear combination of the features with optimized weighting coefficients, $\mathbf{w} = \{w_i\}_{i=1}^n$, as a prediction of the label, $\phi^P(t_k) = w_{0,k} + \sum_{i=1}^n w_{i,k} \phi_i^M$. Based on measured features, the entries of \mathbf{w} are optimized for each time step, t_k , reflecting the time-varying correlations in the dephasing process, captured through the power spectrum.

We demonstrate prediction of a qubit's state subject to stochastic dephasing by performing experiments using the ground-state hyperfine states, $|F=0, m_F=0\rangle$ and $|F=1, m_F=0\rangle$, in trapped $^{171}\text{Yb}^+$ ions as a qubit with transition frequency near 12.6 GHz. A coherent superposition of the qubit states in the measurement basis induced by microwave control³⁵ evolves freely under the influence of an engineered dephasing interaction larger than any intrinsic noise in our experimental system

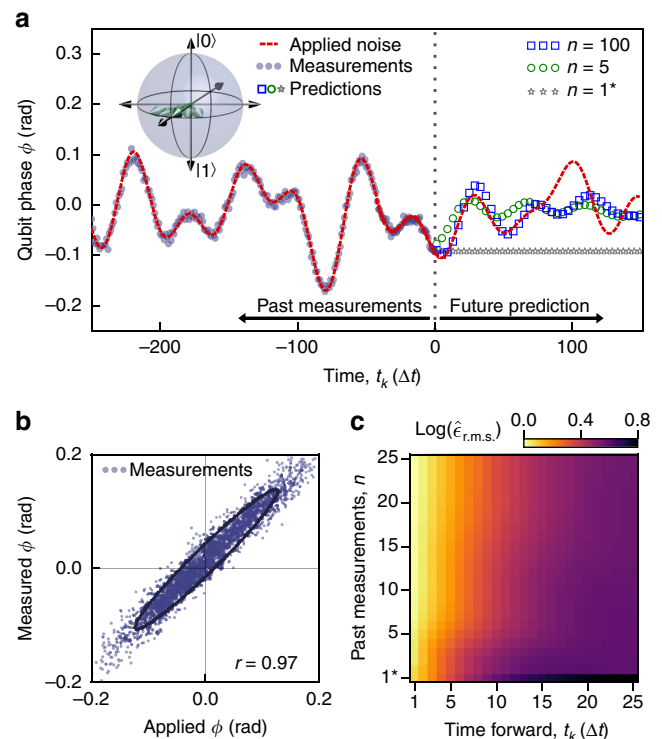


Figure 1 | Prediction of the future evolution of a qubit subject to stochastic dephasing. (a) The average, random phase evolution of the qubit during each measurement under the influence of an engineered noise trace, ϕ^A , is probed via Ramsey spectroscopy and a projective measurement performed before the qubit state is reinitialised and the process repeated (Supplementary Methods). Time is represented in discrete increments of Δt , approximately corresponding to the measurement time. Values of $t_k \leq 0$ refer to past measurements used to make predictions and $t_k > 0$ refer to future predictions. Noise possesses a quasi-white power spectral density up to frequency cut-off ω_c , which we sample at $\omega_s = 40\omega_c$. Future qubit evolution is calculated offline based on these measurements. Data labelled ' $n=1^*$ ' correspond to traditional feedback (no prediction). (inset) Bloch sphere representation of randomization of qubit phase. (b) Correlation between $\phi^M(t_k)$ and $\phi^A(t_k)$ represented as a scatter plot for all measurements in this data set. Ellipses are guides to the eye calculated to have major and minor axis determined by the eigenvectors of the data's covariance matrix. The Pearson product-moment correlation coefficient, r , is calculated to quantify the quality of the measurements—here 97%. (c) Normalized r.m.s. errors $\hat{\epsilon}_{r.m.s.}$ between $\phi^A(t_k)$ and $\phi^P(t_k)$ as a function of past measurements and discrete steps forward in time, averaged over all elements of the data set. Data are normalized to the lowest overall value in the field and are presented using a logarithmic scale to highlight differences over a broad dynamic range. The first row corresponds to traditional feedback.

(Supplementary Methods). In general we work in a regime where the noise evolves slowly during a single measurement period T_M , but we allow the rate at which measurements of qubit-phase evolution are taken—the sampling frequency ω_s —to vary relative to the highest frequency in the noise power spectrum, ω_c (c.f. Fig. 3f). The dephasing noise processes presented here are all derived from a flat-top frequency power spectrum with characteristic cut-off at ω_c . More complex spectra are discussed in Supplementary Discussion and demonstrate similar performance.

An important aspect of our approach is that measurements providing data serving as features may be performed through any suitable protocol. For instance, performing a series of p projective

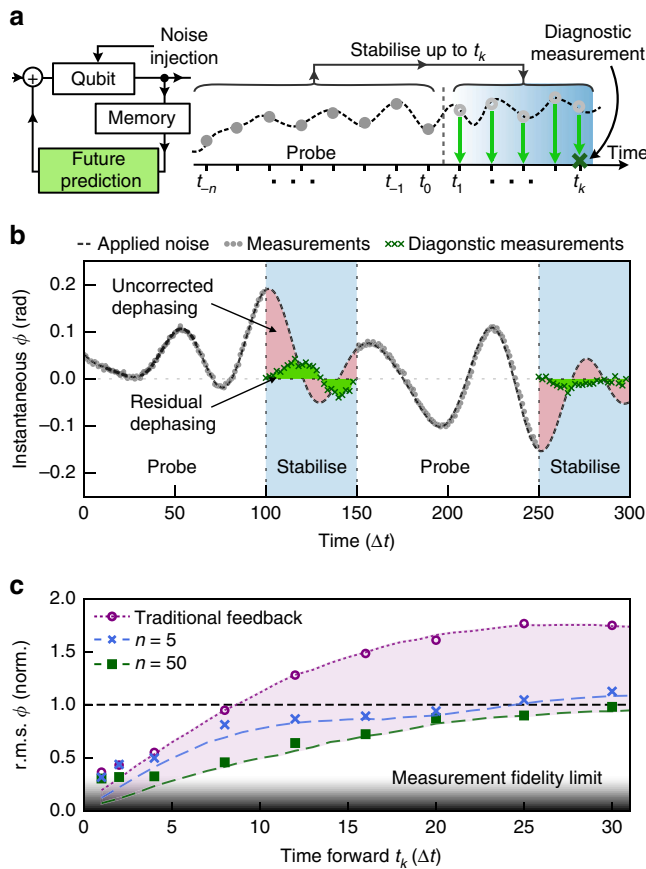


Figure 2 | Experimental time-division multiplexing for qubit stabilization against dephasing. (a) Schematics showing the key aspects of the implementation. Noise is continuously injected into the system. Measurements are taken up to t_0 and processed in real-time to predict future evolution of the qubit phase many time steps ahead, $\phi^P(t_k)$. From t_0 measurement-free compensation based on $\phi^P(t_k)$ is applied during each discrete time step (light green arrows) up to t_k when a diagnostic measurement is performed to verify the accuracy of the prediction/correction process. Full details appear in Supplementary Methods. (b) Probe and stabilization cycles of a time-division multiplexed measurement using $n=100$ past measurements and prediction/correction up to $k=50$ time steps ahead. Green shading indicates reduced residual phase errors. (c) r.m.s. results from time-division multiplexed measurements for different t_k and n compared against traditional feedback and averaged over 50 unique stabilization periods. Data are normalized to the r.m.s. of ϕ^A indicated by the horizontal dashed line. The other dotted/dashed lines are simulations. Markers represent the averaged results of diagnostic measurements. For these data $\omega_s = 40\omega_c$.

measurements on a single qubit to obtain ensemble-averaged information simply sets the scale of the measurement period, $T_M \rightarrow pT_M^{(1)}$, with $T_M^{(1)}$ the duration of a single experiment. Here, we employ a projective measurement that captures statistical information through a spatial ensemble. The impact of such differences is explicitly captured in the sampling frequency of the measurement process.

Forward prediction of stochastic qubit-phase evolution. We begin by accumulating a series of projective measurements of the qubit's phase under engineered dephasing. These serve as training data for the algorithm to optimize the coefficients in \mathbf{w} . We then perform another series of measurements (shown, Fig. 1a) under application of a different noise process possessing similar

statistical characteristics as used in acquiring the training data. This approach ensures that our estimates of prediction accuracy are conservative and exhibit reasonable model robustness and generality. Performing the learning algorithm on a single data set can enhance performance of the prediction algorithm but introduces extreme sensitivity to the input model, ultimately reducing prediction efficacy in the presence of variations in the detailed form of the noise.

An example engineered noise trace in time with overlaid measurement outcomes, ϕ^M , is depicted in Fig. 1a, with 97% correlation between ϕ^M and the applied phase ϕ^A (Fig. 1b). Beyond time t_0 we predict future labels of qubit-phase evolution $\phi^P(t_k)$, up to step t_{150} using a variable number, n , of past measurements and the trained coefficients in \mathbf{w} . Calculated predictions approximate ϕ^A well, reproducing key features including inflection points, maxima and minima as a function of t_k . Our knowledge of the noise is used exclusively for quantitative evaluation of prediction efficacy—it does not enter into the machine-learning algorithm in any form.

Prediction accuracy increases with n , as the algorithm learns more about the temporal correlations in ϕ^A . For values of $k \gtrsim n$, corresponding to prediction times exceeding the range over which the algorithm possesses knowledge about the noise features, the prediction quality diminishes. In addition, over very large values of t_k the prediction tends towards the mean of the noise. Comparing predictive estimation to a 'traditional feedback' model, in which future estimates are based simply on the last measured value $\phi^M(t_0)$, the algorithm shows a distinct advantage as it allows for temporal evolution of the noise in the future.

The quantitative benefits of predictive estimation relative to traditional feedback, and the large t_k behaviour of the predictive algorithm are succinctly captured in the root-mean-square (r.m.s.) prediction error averaged over the entire data set, $\epsilon_{r.m.s.}$, and calculated as a function of t_k and n (Fig. 1c). This demonstrates that even over a large ensemble of predictions the algorithm's advantages remain robust. We now move on to provide examples of real-time qubit stabilization in which the incorporation of future state prediction shows significant advantages over existing techniques.

Time-division multiplexed decoherence suppression. As described above, a reliance on feedback involving frequent projective measurements renders a qubit effectively useless for quantum information or other applications, but omission of stabilization techniques in the presence of dephasing noise may lead to phase errors and eventually to total decoherence. To mitigate the effect of dephasing, we tailor an approach in which we temporally multiplex the necessary measurement and actuation operations in distinct probe and stabilization periods respectively (Fig. 2a,b). During the probe period, a fixed number of measurements are taken and processed in real time. From these measurement outcomes the algorithm produces a prediction of the future time-dependent evolution of the noise during the subsequent stabilization period up to some t_k ; the qubit is dedicated exclusively to measurement of the dephasing process in the probe period. During the stabilization period, corrections are applied during each discrete time step to compensate the predicted stochastic phase evolution, but no measurements are conducted; this permits periods of unsupervised evolution during which the qubit is useful and stabilised against dephasing.

As an example we set the objective of maintaining zero net qubit-phase accumulation (in the rotating frame) during each time step of the stabilization period such that arbitrary high-fidelity operations may be conducted on the qubit; here we apply only the identity. Diagnostic measurements are performed after a




Figure 3 | Experimental comparison of long-term stabilization using traditional and predictive feedback. (a) Schematics showing the key aspects of our cyclic feedback implementation using overlapping measurements. (b–d) Demonstration of feedback accuracy for different sampling frequencies ω_s quantified in units of ω_c , presented through correlation plots (c.f. Fig. 1c) for traditional feedback (blue) and prediction (magenta). Data presented are derived from Fig. 1a. (e) Measured sample variance for various protocols as a function of the number of cycles. Data are normalized to the sample variance of the uncorrected (free-running) signal at 1,000 samples. Each line represents data taken for one particular noise realization and thick lines represent the ensemble average. The inset shows an example suppression of variance over measurement outcomes using predictive against traditional feedback (normalized to the noise amplitude). (f) Sample variance at $N=1,000$ as a function of sampling frequency ω_s in units of ω_c , normalized to the sample variance of the uncorrected signal. The measurement time is fixed and ω_s varied through introduction of dead time between measurements. Dotted lines display simulations and markers the measurement results averaged over ten noise realizations. Error bars represent the s.d. of the mean and the shaded areas show the maximum spread of outcomes. For fixed noise parameters varying ω_s serves as a proxy for changing the ratio of T_M^{-1}/ω_c (Supplementary Discussion). Simulations and measurements in all panels use $n=20$.

variable number of corrections to demonstrate the efficacy of this approach but would not ordinarily be required. Two representative S probe/stabilization cycles are displayed in Fig. 2b showing a reduction in integrated phase error of about 70% after a

stabilization delay of t_{50} during the first cycle and a reduction of about 85% during the second. These improvements are partially limited by measurement fidelity, as illustrated in the ensemble-averaged data (Fig. 2c). Predictive compensation in all tested regimes is superior to corrections based only on traditional feedback down to measurement fidelity limits. Compared against numerical simulations we see that for small t_k the algorithm can provide large relative gains.

Predictive estimation inside a periodic feedback loop. In a second application we employ real-time predictive control in a metrological context. Qubits realised in atoms are frequently used as stable references against which local oscillators (LOs) may be disciplined³⁶. However, stochastic evolution of the LO frequency between interrogations leads to imperfect corrections in the feedback loop. This scenario is commonly encountered when classical processing, actuation and system reinitialisation introduce dead time, producing an effective lag in the feedback loop which degrades the long-term stability of the locked oscillator³⁷. The impact of rapid fluctuations in the LO frequency relative to dead time is generally referred to as the Dick effect³⁸, and represents a significant limiting phenomenon in passive frequency standards using atomic references. The correspondence between LO-induced instabilities in frequency references and dephasing in qubits³⁹ thus invites the application of predictive control in a setting where periodic interrogation and projective measurement are native to the feedback loops used in precision frequency metrology.

The usefulness of predictive estimation in improving correction accuracy inside a feedback loop is demonstrated in Fig. 3b–d, where we plot the predicted phase $\phi^P(t_k)$ (based on two different techniques) against the applied phase error $\phi^A(t_k)$. A prediction with unity correlation to the applied noise would form a diagonal line along $\phi^P = \phi^A$ (similar to Fig. 1b), while imperfect predictions—hence imperfect corrections—result in a dispersion of points around this line in an ellipse.

We vary the sampling frequencies ω_s as a proxy for introducing a variable dead time in the feedback loop (Supplementary Discussion). In a regime where the LO-induced dephasing process evolves slowly, quantified as $\omega_s \gg \omega_c$, both $\phi^M(t_0)$ and the predicted phase $\phi^P(t_k)$ show positive correlation to $\phi^A(t_k)$ (Fig. 3b). As we decrease ω_s , noise evolution during the dead time leads to diminishing correlation between the prediction and actual noise, causing the ellipses to rotate and broaden—a manifestation of the Dick effect.

Predictive estimates are compared with the traditional feedback model described above. For ω_s approaching the Nyquist limit we observe that the traditional prediction can become anticorrelated with the rapidly evolving applied noise (blue ellipse, Fig. 3d), which in real-world applications would lead to an unstable system under feedback. By contrast, using optimized predictions, the decrease in correlation is much slower and the machine-learning algorithm prevents the prediction from ever becoming anticorrelated with the applied dephasing noise. In circumstances tested we always find the optimal prediction correlation $r_P > r_T$ for traditional feedback. Corrections used to discipline the qubit or LO based on predictive estimation can therefore possess enhanced average accuracy relative to traditional feedback.

We now implement real-time evaluation of $\phi^P(t_k)$ inside a feedback loop, demonstrating the ability to improve the individual corrections and ultimately achieve improved long-term stability of the locked qubit. In our experiment we set $n=20$, calculate $\phi^P(t_k)$ on the fly, and cyclically correct based on these predictions (Fig. 3a), again comparing against traditional feedback. The long-term stability achieved under both methods is

calculated via the sample variance⁴⁰ over a variable number of feedback cycles (Fig. 3e).

Over the range of dead times explored experimentally, the use of optimized predictive feedback, in which future estimates are updated as new measurements are acquired in real time, yields net enhancements over the free-running LO (Fig. 3e,f). This includes regimes near the Nyquist limit where rapid evolution of the noise can result in feedback-induced instability as in Fig. 3d. Over most of this range and for the noise parameters we have employed, performance gains over traditional feedback are $\sim 2 \times$ using optimized predictive feedback—a metrologically significant improvement using only enhanced software in the stabilization. Similar performance enhancements have been observed for a wide range of noise spectra and parameters (Supplementary Discussion).

Predictive estimation applied to intrinsic system noise. Finally, with quantitative evaluation of these techniques in hand using engineered noise, we move on to a study of the intrinsic dephasing noise in our system, which arises due to a combination of LO phase noise and magnetic field fluctuations. We perform thousands of sequential projective measurements on the free-running qubit–LO system and process predictions offline. The spectrum of measured fluctuations combines a $1/f^2$ type low-frequency tail with an approximately white plateau, resulting in significant spectral weight near the measurement cycle time. We perform an analysis similar to that presented in Fig. 1, with prediction accuracy quantified using the r.m.s. error between predictions and the future measurement outcomes as a function of t_k (Fig. 4a).

Our machine-learning algorithm enhances the prediction of future qubit evolution by $\sim 30\%$ relative to the r.m.s. error of the uncorrected measurements. We achieve similar performance gains relative to both traditional feedback and the free-running system in calculated sample variance over thousands of correction cycles based on predicted qubit phase, Fig. 4b. In this case the rapid evolution of the noise causes traditional feedback to produce a larger sample variance than free evolution—a situation similar to that experienced in Fig. 3d. The calculated performance enhancements of our method on the intrinsic system noise are significant and show that our algorithm possesses the capability

to improve the stability against the noise background in our system.

Discussion

In this work we have demonstrated the ability to deploy machine-learning techniques to predict and pre-emptively compensate for stochastic qubit dephasing. By exploiting temporal correlations in noise processes, we are able to suppress dephasing during periods when probing the qubit state is not possible, even though we have no deterministic model of the qubit's evolution. Implementing this approach requires neither additional quantum resources nor extra experimental hardware. Instead we rely on software-based machine-learning techniques, which extract optimal performance from information that would have already been collected during common experimental implementations. It has been shown numerically that it is possible to implement an analytical solution to maximally exploit noise correlations captured through the noise power spectrum⁴¹. However in our experimental demonstration the ease of implementation lends itself to use for large values of k and n where prediction is extended far into the future and the computational requirement of large matrix inversions make analytic techniques impractical. In addition, deviations from the idealization of noise characteristics represented by use of a simple power spectral density, as well as correlations appearing in the measurement process, are easily captured by the machine-learning algorithm but invisible to such analytic approaches.

The capability to suppress errors in quantum systems undergoing stochastic evolution has direct implications for the metrology and quantum information communities. In particular the ability to suppress the magnitude of residual dephasing errors makes this technique an attractive complement to open-loop dynamic error suppression for quantum information. Any reduction in the strength of the effective noise experienced by the qubit exponentially improves the fidelity of an operation implemented using dynamic error suppression²⁰. Even in the limit of quasi-static noise, reducing the magnitude of the dephasing error experienced during a dynamically protected operation will improve the ultimate fidelity achievable in a nontrivial quantum logic operation⁴². The complementarity between open- and closed-loop stabilization is a common

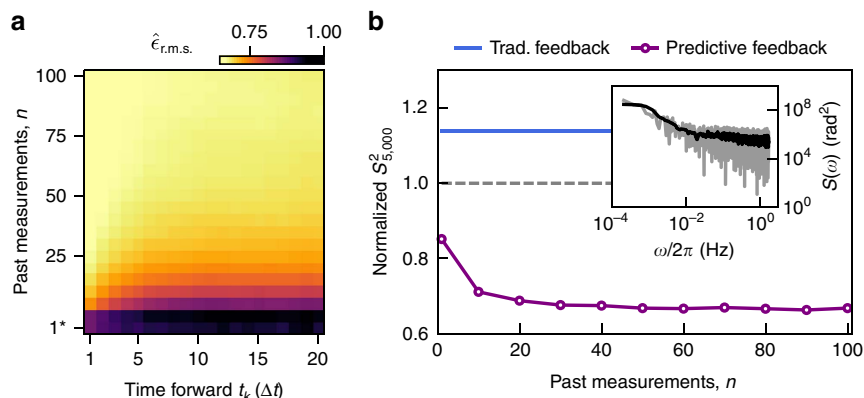


Figure 4 | Application of predictive qubit state estimation to intrinsic system noise. (a) r.m.s. errors between predictions, ϕ^P and actual values ϕ^A for various numbers of past measurements and discrete steps forward in time, averaged over the whole set of validation data. The r.m.s. values are normalized to the r.m.s.d. of the uncorrected data from zero. The bottom row (1*) corresponds to traditional feedback. (b) Sample variance of the corrected measurements averaged over 5,000 cycles, as a function of past measurements used for prediction, normalized to the sample variance of the uncorrected system. The expected sample variance obtained by performing traditional feedback is added for comparison. Data are split into two subsets, where the first 70% serve for training purposes and the remaining 30% are used for validation. (Inset) Power spectrum of a series of projective measurements on the free-running qubit–LO system. The data is overlaid with a smoothed version to visualize the general trend. The maximum frequency in the spectrum corresponds to our sampling frequency and is about 1.7 Hz.

theme in control engineering and translates well to the current setting. Future experiments will involve an expansion to a greater variety of machine-learning algorithms for system characterization and stabilization, and treatment of more complex control scenarios with non-commuting noise terms in the qubit Hamiltonian, nonlinearities in the control, and use of various measurement bases.

Data availability. Data published in this article and the computer code used for simulation is available from the authors.

References

- Gustavson, T. L., Bouyer, P. & Kasevich, M. A. Precision rotation measurements with an atom interferometer gyroscope. *Phys. Rev. Lett.* **78**, 2046–2049 (1997).
- Chou, C. W., Hume, D. B., Rosenband, T. & Wineland, D. J. Optical clocks and relativity. *Science* **329**, 1630–1633 (2010).
- Shah, V., Knappe, S., Schwindt, P. D. D. & Kitching, J. Subpicotesla atomic magnetometry with a microfabricated vapour cell. *Nat. Photon.* **1**, 649–652 (2007).
- Aasi, J. *et al.* Enhanced sensitivity of the LIGO gravitational wave detector by using squeezed states of light. *Nat. Photon.* **7**, 613–619 (2013).
- Wasilewski, W. *et al.* Quantum noise limited and entanglement-assisted magnetometry. *Phys. Rev. Lett.* **104**, 133601 (2010).
- Leibfried, D. *et al.* Toward Heisenberg-limited spectroscopy with multiparticle entangled states. *Science* **304**, 1476–1478 (2004).
- Nielsen, M. A. & Chuang, I. L. *Quantum Computation and Quantum Information* (Cambridge University Press, 2000).
- Lloyd, S. Universal quantum simulators. *Science* **273**, 1073–1078 (1996).
- Somaroo, S., Tseng, C. H., Havel, T. F., Laflamme, R. & Cory, D. G. Quantum simulations on a quantum computer. *Phys. Rev. Lett.* **82**, 5381–5384 (1999).
- Somma, R., Ortiz, G., Gubernatis, J. E., Knill, E. & Laflamme, R. Simulating physical phenomena by quantum networks. *Phys. Rev. A* **65**, 042323 (2002).
- Aspuru-Guzik, A. & Walther, P. Photonic quantum simulations. *Nat. Phys.* **8**, 285–291 (2012).
- Bloch, I., Delibard, J. & Nascimbene, S. Quantum simulations with ultracold quantum gases. *Nat. Phys.* **8**, 267–276 (2012).
- Blatt, R. & Roos, C. F. Quantum simulations with trapped ions. *Nat. Phys.* **8**, 277–284 (2012).
- Stengel, R. F. *Optimal Control and Estimation* (Dover Publications Inc., 1994).
- Clark, J. W., Lucarelli, D. G. & Tarn, T. J. Control of quantum systems. *Int. J. Mod. Phys. B* **17**, 5397–5411 (2003).
- Wiseman, H. M. & Doherty, A. C. Optimal unravellings for feedback control in linear quantum systems. *Phys. Rev. Lett.* **94**, 1–4 (2005).
- Nurdin, H. I., James, M. R. & Petersen, I. R. Coherent quantum LGQ control. *Automatica* **45**, 1837–1846 (2009).
- Biercuk, M. J., Doherty, A. C. & Uys, H. Dynamical decoupling sequence construction as a filter-design problem. *J. Phys. B* **44**, 154002 (2011).
- Viola, L. & Lloyd, S. Dynamical suppression of decoherence in two-state quantum systems. *Phys. Rev. A* **58**, 2733–2744 (1998).
- Soare, A. *et al.* Experimental noise filtering by quantum control. *Nat. Phys.* **10**, 825–829 (2014).
- Bluhm, H., Foletti, S., Mahalu, D., Umansky, V. & Yacoby, A. Enhancing the coherence of a spin qubit by operating it as a feedback loop that controls its nuclear spin bath. *Phys. Rev. Lett.* **105**, 216803 (2010).
- Vijay, R. *et al.* Stabilizing rabi oscillations in a superconducting qubit using quantum feedback. *Nature* **490**, 77–80 (2012).
- Shulman, M. D. *et al.* Suppressing qubit dephasing using real-time Hamiltonian estimation. *Nat. Commun.* **5**, 5156 (2014).
- Steck, D. A., Jacobs, K., Mabuchi, H., Bhattacharya, T. & Habib, S. Quantum feedback control of atomic motion in an optical cavity. *Phys. Rev. Lett.* **92**, 223004 (2004).
- Wilson, D. J. *et al.* Measurement-based control of a mechanical oscillator at its thermal decoherence rate. *Nature* **524**, 325–329 (2015).
- Wieczorek, W. *et al.* Optimal state estimation for cavity optomechanical systems. *Phys. Rev. Lett.* **114**, 223601 (2015).
- Sayrin, C. *et al.* Real-time quantum feedback prepares and stabilizes photon number states. *Nature* **477**, 73–77 (2011).
- Gillett, G. G. *et al.* Experimental feedback control of quantum systems using weak measurements. *Phys. Rev. Lett.* **104**, 2–5 (2010).
- Hirose, M. & Cappellaro, P. Coherent feedback control of a single qubit in diamond. *Nature* **532**, 77–80 (2016).
- Landau, I. D., Lozano, R., M'Saad, M. & Karimi, A. (ed.). *Adaptive Control, Algorithms, Analysis and Applications* (Springer, 2011).
- Nielsen, A. E. B., Hopkins, A. S. & Mabuchi, H. Quantum filter reduction for measurement-feedback control via un-supervised manifold learning. *New J. Phys.* **11**, 105043 (2009).
- Tan, D., Weber, S. J., Siddiqi, I., Mølmer, K. & Murch, K. W. Prediction and retrodiction for a continuously monitored superconducting qubit. *Phys. Rev. Lett.* **114**, 090403 (2015).
- Murphy, K. P. *Machine Learning, A Probabilistic Approach* (MIT Press, 2012).
- Van Harlingen, D. J., Plourde, B. L. T., Robertson, T. L., Reichardt, P. A. & Clarke, J. in *Quantum Computing and Quantum Bits in Mesoscopic Systems* (eds Leggett, A., Ruggiero, B., & Silvestini, P.) 171–184 (Kluwer Academic Press, 2004).
- Soare, A. *et al.* Experimental bath engineering for quantitative studies of quantum control. *Phys. Rev. A* **89**, 042329 (2014).
- Audoin, C. & Guinot, B. *The Measurement of Time* 1st English translation edn (Cambridge University Press, 2001).
- Greenhall, C. A. A derivation of the long-term degradation of a pulsed atomic frequency standard from a control-loop model. *IEEE Trans. Ultrason. Ferroelect. Freq. Control* **45**, 895–898 (1998).
- Dick, G. Local oscillator induced instabilities in trapped ion frequency standards. In *Proceedings of the 19th Annual Precise Time and Time Interval (PTTI) Systems and Applications and Planning Meeting* (1987).
- Ball, H., Oliver, W. D. & Biercuk, M. J. The role of master clock stability in quantum information processing. *Nature Quantum Information* **2**, 16033 (2016).
- Rutman, J. Characterization of phase and frequency instabilities in precision frequency sources: fifteen years of progress. *Proc. IEEE* **66**, 1048–1075 (1978).
- Sastrawan, J., Jones, C., Akhalwaya, I., Uys, H. & Biercuk, M. J. Analytically exploiting noise correlations inside the feedback loop to improve locked-oscillator performance. *Phys. Rev. E* **94**, 022204 (2016).
- Kabytayev, C. *et al.* Robustness of composite pulses to time-dependent control noise. *Phys. Rev. A* **90**, 012316 (2014).

Acknowledgements

The authors acknowledge D. Hayes, M.C. Jarratt and A. Soare, for contributions to the experimental system and C. Edmunds, C. Ferrie and C. Hempel for useful discussions. Work partially supported by the ARC Centre of Excellence for Engineered Quantum Systems CE110001013, ARC Discovery Project DP130103823, the Intelligence Advanced Research Projects Activity (IARPA) through the ARO, the US Army Research Office under Contract W911NF-12-R-0012, and a private grant from H. & A. Harley

Author contributions

S.M. and V.F. developed experimental hardware and the experimental control system, performed numerical simulations, obtained the presented data, and carried out the data analysis with guidance from M.J.B. M.J.B. conceived the experiment and led development of the experimental system. J.S. and S.D. obtained preliminary data. S.M., V.F. and M.J.B. wrote the manuscript.

Additional information

Supplementary Information accompanies this paper at <http://www.nature.com/naturecommunications>

Competing financial interests: The authors declare no competing financial interests.

Reprints and permission information is available online at <http://npg.nature.com/reprintsandpermissions/>

How to cite this article: Mavadia, S. *et al.* Prediction and real-time compensation of qubit decoherence via machine learning. *Nat. Commun.* **8**, 14106 doi: 10.1038/ncomms14106 (2017).

Publisher's note: Springer Nature remains neutral with regard to jurisdictional claims in published maps and institutional affiliations.



This work is licensed under a Creative Commons Attribution 4.0 International License. The images or other third party material in this article are included in the article's Creative Commons license, unless indicated otherwise in the credit line; if the material is not included under the Creative Commons license, users will need to obtain permission from the license holder to reproduce the material. To view a copy of this license, visit <http://creativecommons.org/licenses/by/4.0/>

© The Author(s) 2017